

Claims

1. A data switching device having a plurality of ingress routers, a plurality of egress routers, a switching matrix and a connection controller,

the switching matrix having input ports connected to respective said
 5 ingress routers and output ports connected to respective said egress routers, and controlled by the controller to form connections between pairs of input and output ports;

each ingress router including one or more virtual output queues for each egress router, each virtual output queue being arranged to store fixed
 10 length cells having a header defining the egress router to be used in the switching matrix connection;

each ingress router being arranged, upon receipt of a new cell by that ingress router, to store the cell in a said virtual output queue of the ingress router corresponding to the egress router for the cell;

15 characterized in that:

each input port of the switching matrix includes for each virtual output queue in the ingress router connected to that input port a respective head of queue buffer;

the switching matrix is arranged, upon the switching matrix forming a
 20 connection between a given input port and output port, to transmit to that output port a cell from a corresponding one of the head of queue buffers;

each egress router is arranged, upon error free receipt by that egress router of a cell from one of the virtual output queues of one of the ingress routers, to transmit a receipt signal to that ingress router;

25 each ingress router is arranged, upon storing a cell in a said virtual output queue, and, if a credit count, indicative of the number of free cells of the corresponding head of queue buffer, is not zero, to transmit a replication

005090" 00542860

of the cell to that head of queue buffer and a connection request to the controller; and

each ingress router is further arranged to retain each received cell in the corresponding virtual output queue until receiving the corresponding receipt signal.

2. A device according to claim 1 in which, upon the controller causing a connection to switch a cell of an ingress router through the switching matrix, a connection grant signal is transmitted to that ingress router, and increments the credit count by one.

10 3. A device according to claim 1 in which, upon said replication of the cell to the head of queue buffer, the respective credit count is decremented by one.

4. A device according to claim 1 in which the virtual output queues are segregated into two areas, a first area containing cells waiting for replication to the corresponding head of queue buffer, and a second area containing cells replicated to the head of queue buffer, and, upon determining that there is at least one cell in the first area and that the number of free cells of the corresponding head of queue buffer is not zero, a replication of at least one cell in the first area is transmitted to that head of queue buffer, the cell is transferred to the second area, and a connection request is transmitted to the controller.

5. A device according to claim 1 in which each cell is associated with a priority level, said virtual output queues comprising a virtual output queue for cells of each respective priority level, said controller determining, in tandem with which pairs of input and output ports will be connected, the priority level of the cell to be transmitted between them.

6. A device according to claim 1 in which the controller determines whether any given one of the cells in the virtual output queues can be transmitted between the appropriate pair of input and output ports without preventing the transmission of a cell in a virtual output queue between another pair of input and output ports, and in this case causes that given cell to be transmitted.

7. A device according to claim 1 in which each egress router is arranged to detect that a cell transmitted by the switching matrix has not been received correctly, and in this case transmits a re-transmission request to the corresponding ingress router.

8. A device according to claim 4 in which, upon receiving the re-transmission request, the ingress router transfers the corresponding cell in the second area into the first area, and transmits a corresponding connection request to the controller.

9. A method of operating a data switching device having a plurality of ingress routers, a plurality of egress routers, a switching matrix and a connection controller, and

the switching matrix having input ports connected to respective said ingress routers and output ports connected to respective said egress routers, and controlled by the controller to form connections between pairs of input and output ports;

the method comprising the step of:

maintaining at each ingress router one or more virtual output queues for each egress router, each virtual output queue being arranged to store fixed length cells having a header defining the egress router to be used in the switching matrix connection, and each ingress router being arranged upon receipt of a new cell to store the cell in a said virtual output queue of the ingress router corresponding to the egress router for the cell;

characterized in that the method further includes the steps of:

maintaining at each input port of the switching matrix for each virtual output queue in the ingress router connected to that input port, a respective head of queue buffer;

5 upon the switching matrix forming a connection between a given input port and output port, the switching matrix transmitting from that input port to that output port a cell from one of the one or more corresponding head of queue buffers,

10 upon error free receipt by an egress router of a cell from one of the virtual output queues of one of the ingress routers, that egress router transmitting a receipt signal to that ingress router,

15 upon receipt of a new cell by one of the ingress routers, and, if a credit count, indicative of the number of free cells of the corresponding head of queue buffer, is not zero, the ingress router transmitting a replication of the cell to that head of queue buffer and a connection request to the controller; and

each ingress router retaining each received cell in the corresponding virtual output queue until receiving the corresponding receipt signal.

20 10. A method according to claim 9 in which, upon the controller causing a connection to switch a cell of an ingress router through the switching matrix (57; 66), a connection grant signal is transmitted to that ingress router, and increments the credit count by one.

25 11. A method according to claim 9 in which, upon said replication of the cell to the head of queue buffer, the respective credit count is decremented by one.

12. A method according to claim 9 in which the virtual output queues are segregated into two areas, a first area containing cells waiting for replication to the corresponding head of queue buffer, and a second area containing cells replicated to the head of queue buffer, and, upon determining that there is at least one cell in the first area and that the number of free cells of the corresponding head of queue buffer is not zero, a replication of at least one cell in the first area is transmitted to that head of queue buffer, the cell is transferred to the second area, and a connection request is transmitted to the controller.

13. A method according to claim 9 in which each cell is associated with a priority level, said virtual output queues comprising a virtual output queue for cells of each respective priority level, said controller determining, in tandem with which pairs of input and output ports will be connected, the priority level of the cell to be transmitted between them.

14. A method according to claim 9 in which the controller determines whether any given one of the cells in the virtual output queues can be transmitted between the appropriate pair of input and output ports without preventing the transmission of a cell in a virtual output queue between another pair of input and output ports, and in this case causes that given cell to be transmitted.

15. A method according to claim 9 in which each egress detects that a cell transmitted by the switching matrix has not been received correctly, and in this case transmits a re-transmission request to the corresponding ingress router.

16. A method according to claim 12 in which, upon receiving the re-transmission request, the ingress router transfers the corresponding cell in the

6